



US009274874B1

(12) **United States Patent**  
**Chamness et al.**

(10) **Patent No.:** **US 9,274,874 B1**  
(45) **Date of Patent:** **Mar. 1, 2016**

(54) **AUTOMATED DEFECT DIAGNOSIS FROM MACHINE DIAGNOSTIC DATA**

(71) Applicant: **EMC Corporation**, Hopkinton, MA (US)

(72) Inventors: **Mark Chamness**, Menlo Park, CA (US);  
**Eric Schnegelberger**, Youngsville, NC (US)

(73) Assignee: **EMC Corporation**, Hopkinton, MA (US)

(\*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 194 days.

(21) Appl. No.: **14/042,491**

(22) Filed: **Sep. 30, 2013**

(51) **Int. Cl.**  
**G06F 11/00** (2006.01)  
**G06F 11/07** (2006.01)

(52) **U.S. Cl.**  
CPC ..... **G06F 11/0775** (2013.01); **G06F 11/0781** (2013.01)

(58) **Field of Classification Search**  
CPC ..... G06F 11/0775; G06F 11/0781; G06F 11/0787; G06F 11/079; G06F 11/0793; G06F 11/2257; G06F 11/2263; G06F 11/3058; G06F 11/3072; G06F 11/3086  
See application file for complete search history.

(56) **References Cited**

**U.S. PATENT DOCUMENTS**

5,928,369 A	7/1999	Keyser et al.	
7,203,864 B2	4/2007	Goin et al.	
7,788,536 B1 *	8/2010	Qureshi et al.	714/38.14
7,937,623 B2 *	5/2011	Ramacher et al.	714/37

2004/0012808 A1	1/2004	Payne et al.	
2005/0114739 A1	5/2005	Gupta et al.	
2005/0120105 A1	6/2005	Popescu et al.	
2005/0278213 A1	12/2005	Faihe	
2006/0085785 A1	4/2006	Garrett	
2007/0220036 A1	9/2007	Das et al.	
2007/0226546 A1	9/2007	Asthana et al.	
2008/0040628 A1	2/2008	Mandal	
2008/0126773 A1	5/2008	Martinez et al.	
2008/0133540 A1 *	6/2008	Hubbard et al.	707/10
2010/0064285 A1	3/2010	Dechovich et al.	
2010/0077008 A1	3/2010	Davis et al.	
2011/0131453 A1	6/2011	Fernandess et al.	
2012/0041575 A1	2/2012	Maeda et al.	
2012/0078837 A1	3/2012	Bagchi et al.	
2012/0130968 A1 *	5/2012	Mule	707/706
2012/0226943 A1	9/2012	Alderman et al.	
2012/0303618 A1	11/2012	Dutta et al.	
2013/0007527 A1 *	1/2013	Petukhov et al.	714/37
2013/0073552 A1	3/2013	Rangwala et al.	
2013/0091266 A1 *	4/2013	Bhave et al.	709/224
2013/0254599 A1	9/2013	Katkar et al.	

**OTHER PUBLICATIONS**

Chamness et al., U.S. Appl. No. 13/928,810, filed Jun. 27, 2013.  
Chamness et al., U.S. Appl. No. 13/928,937, filed Jun. 27, 2013.  
Schnegelberger et al., U.S. Appl. No. 14/042,494, filed Sep. 30, 2013.  
Schnegelberger et al., U.S. Appl. No. 14/042,501, filed Sep. 30, 2013.  
Schnegelberger et al., U.S. Appl. No. 14/038,312, filed Sep. 26, 2013.

\* cited by examiner

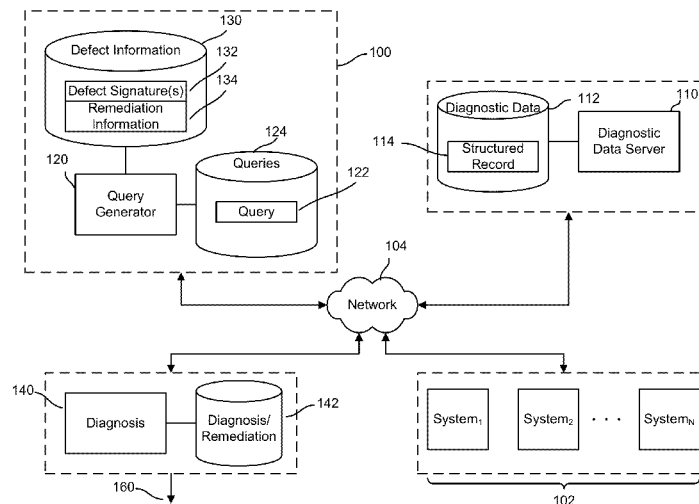
*Primary Examiner* — Joseph Kudirka

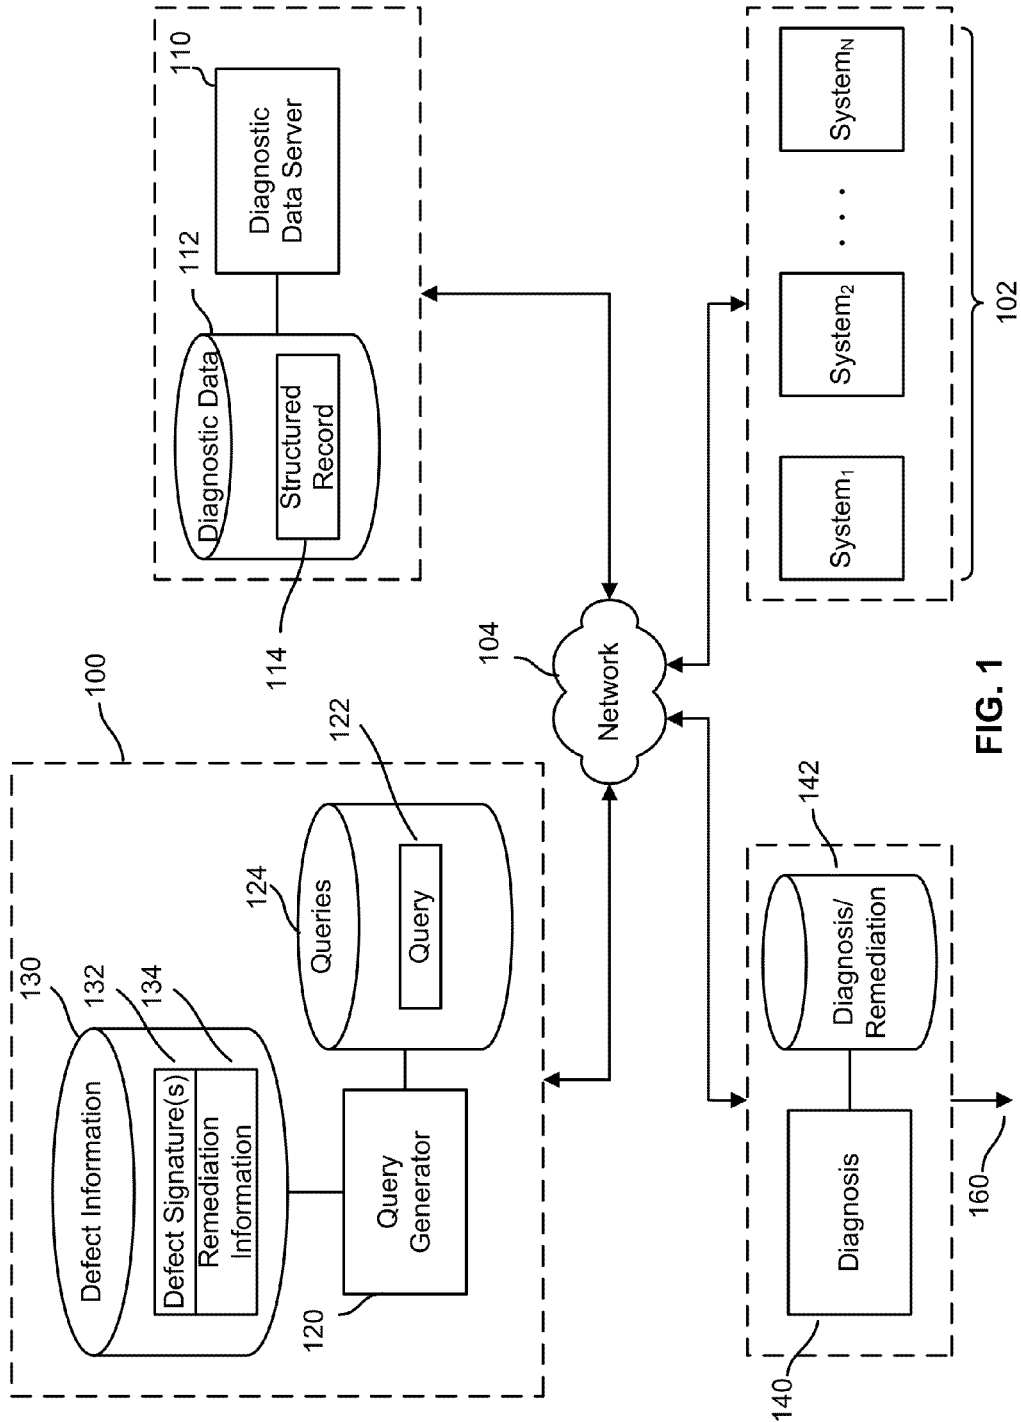
(74) *Attorney, Agent, or Firm* — Van Pelt, Yi & James LLP

(57) **ABSTRACT**

Diagnosis of defect(s) in a system is disclosed. A defect signature-based query is performed against system diagnostic data stored in one or more structured records. It is determined that a defect signature is associated with a system based at least in part on the query. Remediation information generated based at least in part on the defect signature and the system diagnostic data may be output.

**20 Claims, 6 Drawing Sheets**





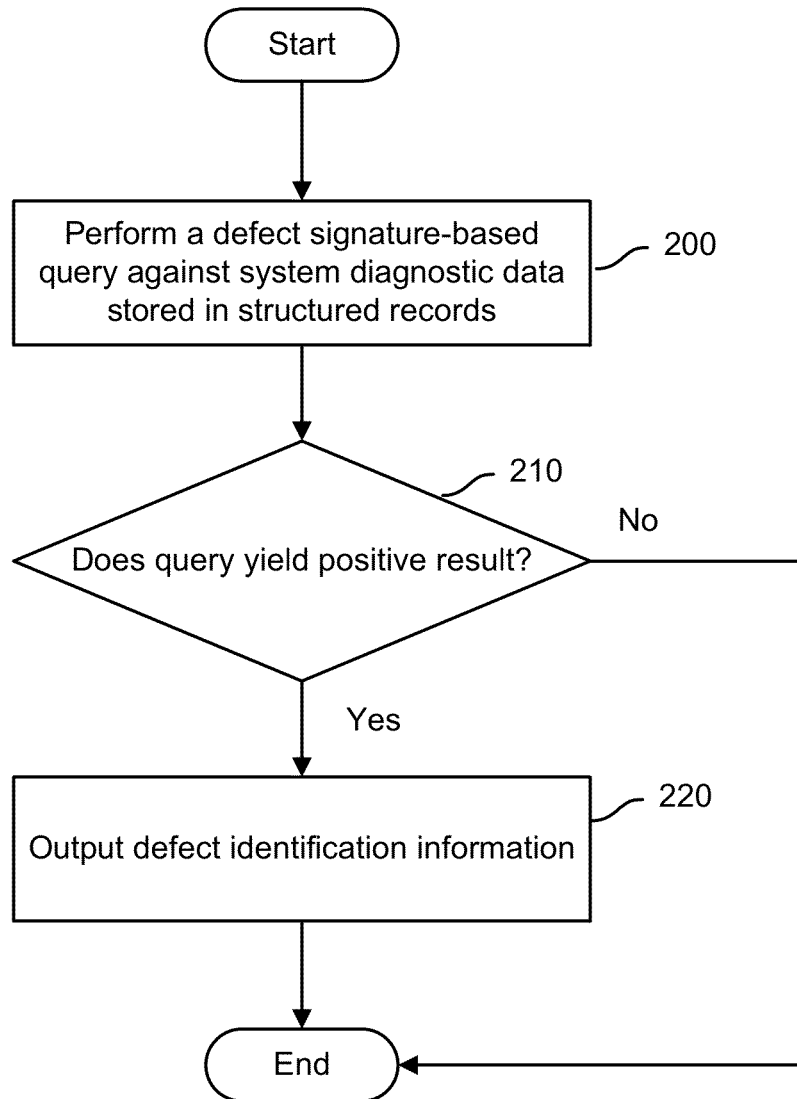
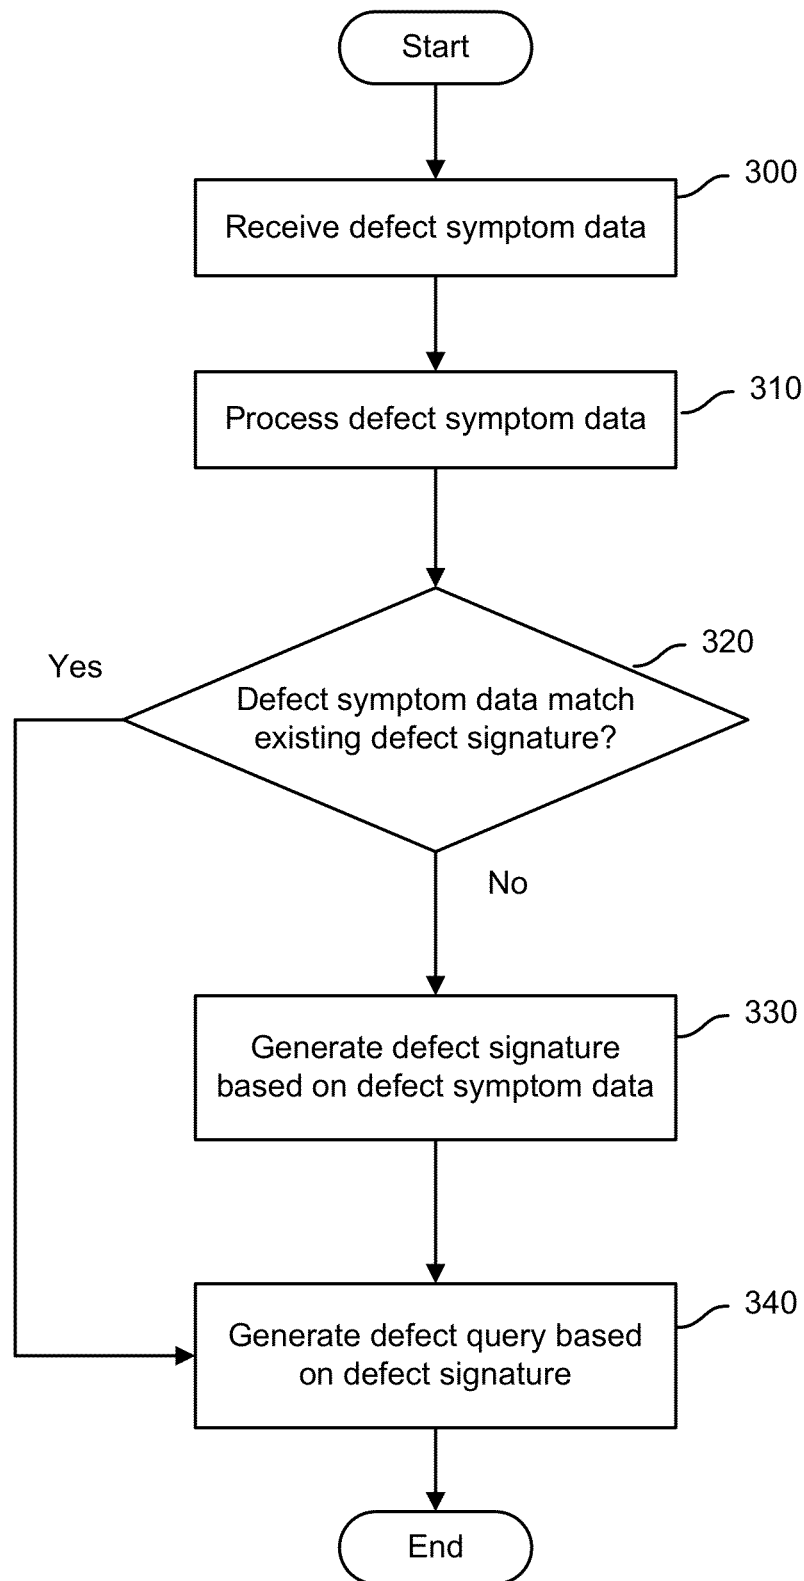
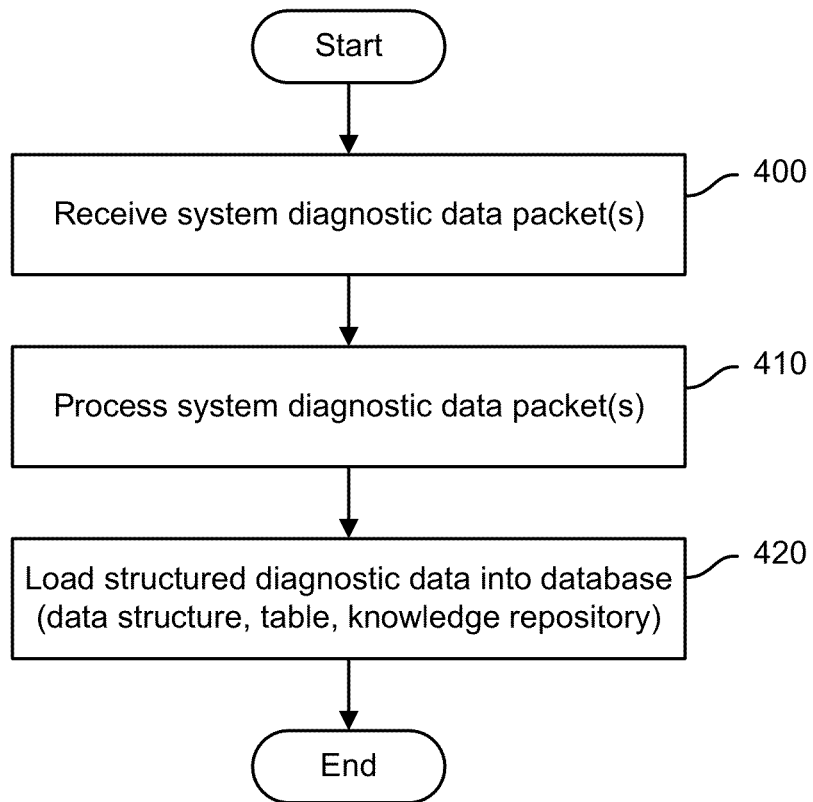


FIG. 2

**FIG. 3**

**FIG. 4**

500

510

Symptom Data:

From: 2013-01-01

To: 2013-02-01

520

Defect (Bug) Filters:

Domain: --All--

Component: --All--

Bug ID: --All--

530

Search:

Serial Number: 123456

Case Number:

540

Scan Symptom Data

FIG. 5

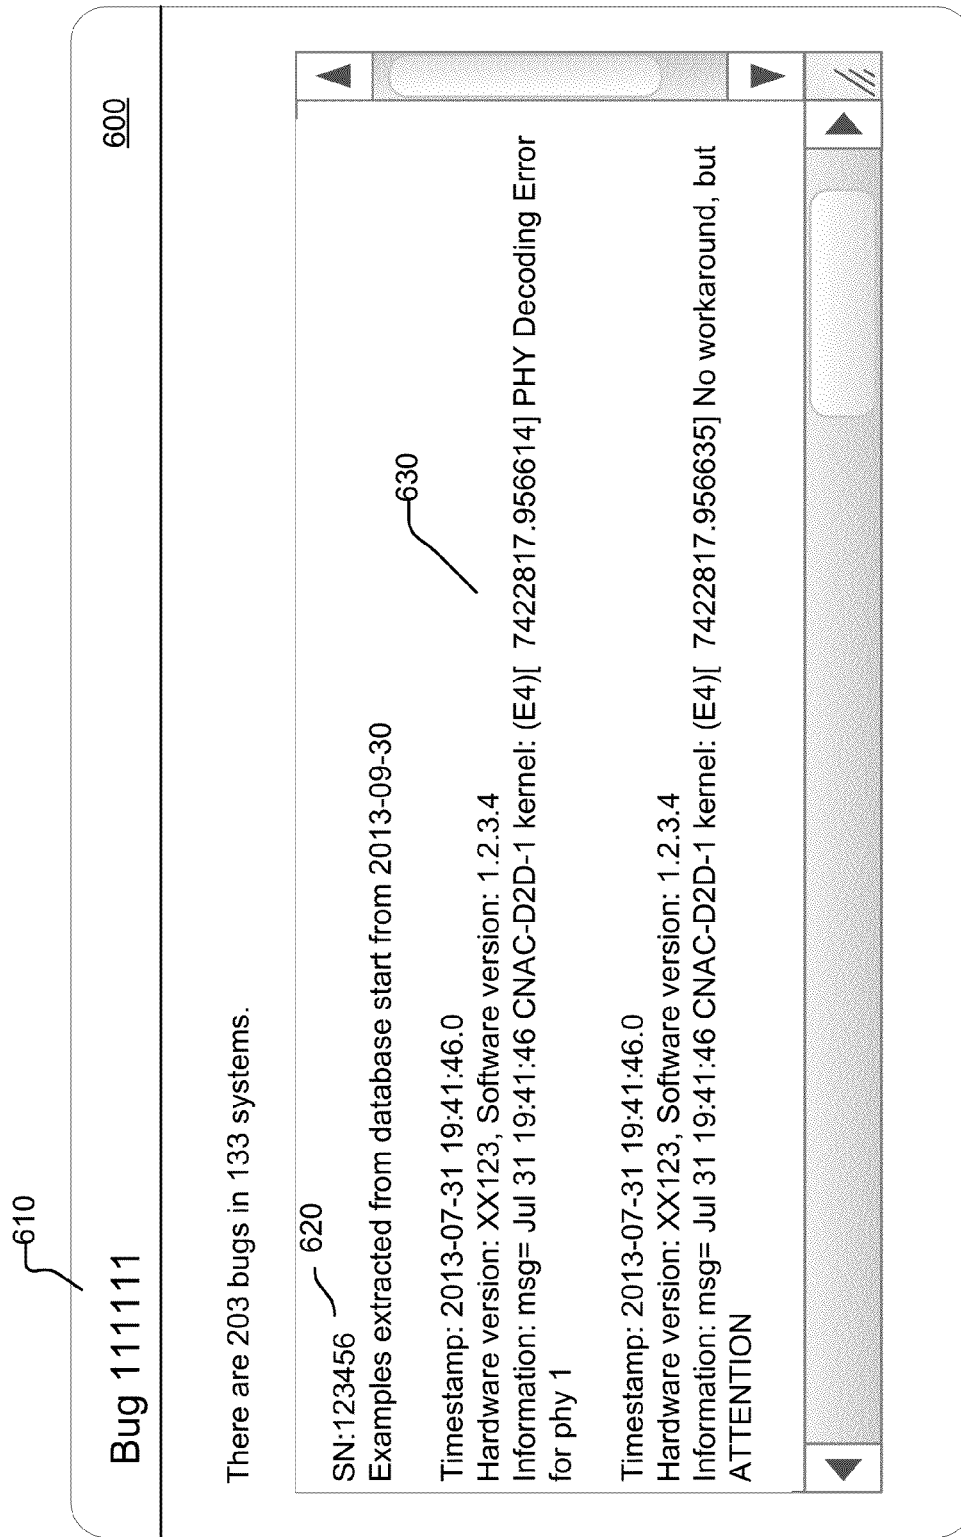


FIG. 6

1

## AUTOMATED DEFECT DIAGNOSIS FROM MACHINE DIAGNOSTIC DATA

### BACKGROUND OF THE INVENTION

Technical support personnel typically spend a great deal of time manually reviewing system data to diagnose and resolve problems associated with a system. For example, a support engineer may receive a call from a customer regarding a defect in the customer's system. Often based on rough description of a defect from the customer, technical support personnel must manually search through system diagnostic data to identify defect symptoms. This process of manually searching through diagnostic data may be inefficient and costly for a technical support department.

### BRIEF DESCRIPTION OF THE DRAWINGS

Various embodiments of the invention are disclosed in the following detailed description and the accompanying drawings.

FIG. 1 is a diagram showing an automated support system for defect diagnosis.

FIG. 2 is a flowchart illustrating embodiments of a process for defect diagnosis.

FIG. 3 is a flowchart illustrating embodiments of a process for generating defect signatures.

FIG. 4 is a flowchart illustrating embodiments of a process for processing system diagnostic data.

FIG. 5 is a diagram depicting an example interface for defect diagnosis.

FIG. 6 is a diagram depicting an example interface for displaying defect diagnosis results.

### DETAILED DESCRIPTION

The invention can be implemented in numerous ways, including as a process; an apparatus; a system; a composition of matter; a computer program product embodied on a computer readable storage medium; and/or a processor, such as a processor configured to execute instructions stored on and/or provided by a memory coupled to the processor. In this specification, these implementations, or any other form that the invention may take, may be referred to as techniques. In general, the order of the steps of disclosed processes may be altered within the scope of the invention. Unless stated otherwise, a component such as a processor or a memory described as being configured to perform a task may be implemented as a general component that is temporarily configured to perform the task at a given time or a specific component that is manufactured to perform the task. As used herein, the term "processor" refers to one or more devices, circuits, and/or processing cores configured to process data, such as computer program instructions.

A detailed description of one or more embodiments of the invention is provided below along with accompanying figures that illustrate the principles of the invention. The invention is described in connection with such embodiments, but the invention is not limited to any embodiment. The scope of the invention is limited only by the claims and the invention encompasses numerous alternatives, modifications and equivalents. Numerous specific details are set forth in the following description in order to provide a thorough understanding of the invention. These details are provided for the purpose of example and the invention may be practiced according to the claims without some or all of these specific details. For the purpose of clarity, technical material that is

2

known in the technical fields related to the invention has not been described in detail so that the invention is not unnecessarily obscured.

Diagnostic methods and systems which automatically identify defects (e.g., bugs, software defects, misconfigurations, etc.) and other characteristics associated with a system (e.g., a serviced system, computing device, storage device) are disclosed herein. Such methods and systems may be useful for technical support personnel to efficiently diagnose defects in a system. In some embodiments, defect signature-based queries may be performed against structured records of system diagnostic data to determine whether a defect is associated with one or more systems. Upon determination that a defect is associated with a system, remediation approaches may be implemented to correct the defect.

FIG. 1 is a diagram showing an automated support system for defect diagnosis. An automated support system for defect diagnosis **100** may be associated with one or more systems **102** (e.g., serviced system(s), data storage system(s), Data Domain deduplication storage systems, back-up storage devices, servers, or any computer-related system). Systems **102** may communicate with system **100** via a network **104** (e.g., local area network, a wide area network, the Internet, or any appropriate communication network).

In some embodiments, system **100** may be associated with an organization's computer network (e.g., an enterprise computer network), and the systems **102** may include one or more devices managed within that organization. In other embodiments, system **100** may be associated with a system manufacturer (e.g., a data storage product manufacturer), and the systems **102** may include devices purchased from and supported by the system manufacturer. The system manufacturer may, for example, provide technical support to ensure optimal function of the systems **102**.

One or more systems **102** (e.g., serviced systems, storage systems, servers, computing devices) may send system diagnostic data to system **100**. The system diagnostic data may, for example, be received at system **100** (e.g., periodically, once a day). Any appropriate protocol may be used to transport the diagnostic data, including syslog, email message(s), Hypertext Transfer Protocol (HTTP), Hypertext Transfer Protocol Secure (HTTPS), or Simple Network Management Protocol (SNMP). In one example, the diagnostic data may include autosupport files (e.g. a document including diagnostic data for a system). Upon receipt, the diagnostic data may be parsed and otherwise processed by the diagnostic data server **110** to generate one or more structured records **114**. The structured records **114** (e.g., structured diagnostic records) may be stored in a database of system diagnostic data **112** that includes diagnostic data associated with one or more systems **102**.

In some embodiments, a query generator **120** may generate one or more defect signature-based queries **122** (e.g., defect-based queries). A query **122** may be generated based on a defect signature **132** (e.g., bug signature, health signature) associated with a defect and other parameters. A defect may include a bug (e.g., an undesirable condition related to a system), error, or other condition associated with a system that causes the system to deviate from a standard/optimal operation. Defect symptom data may include data that indicates the presence of a defect. Thus, the presence of defect symptom data in system diagnostic data may indicate that an associated defect is present and/or has manifested itself in a particular system **102** in which the defect symptom data is found to be present. Defect symptom data may include, for example, environmental measurements, component status, configuration settings, policy definitions, logs of system



events (e.g., system failures), logs of system behavior (e.g., central processing unit (CPU) usage, input/output), and/or other information. A defect definition may include information used to characterize a specific defect. And, a defect signature **132** (e.g., a bug signature, health signature) may include defect symptom data which is associated with (e.g., matches, is representative of) a defect. In some cases, a defect signature **132** may include product(s) (e.g., software, hardware, devices) and code level(s) to which the defect is applicable or in which the defect is known to appear. A bug signature may include symptom data associated with a bug. A health signature may include system data representing normal or predicted operation of a system.

According to some embodiments, defect information **130** may be received (e.g., from an engineer, programmer, or other user) at system **100**. The defect information **130** may be processed (e.g., parsed and otherwise processed) into one or more defect symptom records (e.g., collectively defect symptom data). The processed or structured defect symptom records may be stored in a database (e.g., a table, data structure, array, or index). Defect signatures **132** may be generated based at least in part on one or more defect symptom records (e.g., defect symptom data).

A defect signature-based query **122** may be generated based on one or more defect signatures **132** (e.g., by a query generator **120**). Queries **122** may include functions (e.g., Structured Query Language (SQL) functions, PostgreSQL (PL/pgSQL) functions, or other functions executed in any suitable language) configured to query a given data set (e.g., a table, index, array, or other data structure) based on data associated with a defect signature **132** and, possibly, additional parameters. The data associated with a defect signature **132** may include, for example, defect symptom data (e.g., string(s) identifying defect symptoms), configuration parameters (e.g., software/hardware versions), and other data. Other parameters may include, for example, data associated with a system (e.g., a system identifier, serial number associated with a system), a time period (e.g., within one hour, over the last week, etc.), a frequency of occurrences during a window or period of time (e.g., ten times in one day), a domain and component associated with a system, and other suitable parameters. In one example, a defect signature-based query may include a regular expression (regexp), and defect signatures may be represented as regular expression patterns. In some embodiments, a query **122** may include logical operators (e.g., AND, OR, EITHER, IF, or any other logical operator). Logical operators may be used to, for example, combine defect symptoms, parameters associated with a system (e.g., a system identifier), and/or other parameters to generate a customized query. A query **122** may also include additional function calls (e.g., nested function(s)). Ultimately, a query **122** may be based on any combination of parameters, logical operators, techniques, and/or functions. In various embodiments, a query **122** invokes and makes use of a structure that has been imparted to diagnostic data in the course of generating structured records **114** based on received diagnostic data, to search for the presence of defect symptoms in an efficient and computationally feasible way.

According to some embodiments, a diagnostic engine **140** may be configured to perform queries **122** against system diagnostic data stored in one or more structured records **114**. A diagnostic engine **140** may include an application for diagnosing defects (e.g., a web application for diagnosing defects). To perform a query, diagnostic engine **140** may communicate with a diagnostic data server **110** and query generator **120** via network **104**. Diagnostic engine **140** may perform a defect signature-based query using a query **122**

generated by the query generator **120**. The query **122** may be based at least in part on a defect signature **132**, data associated with a system (e.g., a system serial number), and other information. A query **122** may also be generated based on parameters received from diagnostic engine **140** (e.g., based on user input).

The defect signature-based query **122** may be performed against system diagnostic data stored in one or more structured records **114**. In some embodiments, diagnostic engine **140** may communicate with diagnostic data server **110** to request system diagnostic data (e.g., structured records **114**) for the query.

In some embodiments, it may be determined that a defect signature **132** is associated with a system based at least in part on the defect signature-based query **122**. For example, in the event that the defect signature-based query **122** (e.g., based on defect symptoms, data associated with a system, and/or other parameters) results in a positive result (e.g., a match, hit, or correlation), it may be determined that a defect (e.g., associated with the defect signature) is associated with the system.

Upon a determination that the defect is associated with the system, diagnostic engine **140** and/or a device associated therewith, may output defect identification information **160**. Defect identification information **160** may include details regarding the identification of the defect in the system diagnostic data **114**.

In some embodiments, remediation information **134** may be generated based at least in part on the defect signature **132**, system diagnostic data **112**, and/or other information. Remediation information **134** may include, for example, steps to correct/remediate a defect, recommended software/hardware upgrades to correct the defect, a script to correct the defect (e.g., computer-executable instructions to remedy a defect (bug), a software patch), background information regarding the defect, and/or any other information related to the defect.

According to some embodiments, remediation information **134** may be associated with a defect signature **132**. Remediation information **134**, for example, may include an element of the defect signature **132**, may be linked to a defect signature **132**, or otherwise associated (e.g., via metadata) with a defect signature. In other embodiments, remediation information **134** may be generated based on a determination that the defect signature **132** is associated with a system **102**. The remediation information **134** therefore may be based on the diagnosis or identification of defect symptoms in a particular set of system diagnostic data **112** independent of the defect signature.

FIG. 2 is a flowchart illustrating embodiments of a process for defect diagnosis. In some embodiments, the process is performed by system **100** depicted in FIG. 1. At **200**, a defect signature-based query is performed against system diagnostic data stored in one or more structured records. In various embodiments, a defect signature-based query may include a regular expression (regexp), and defect signatures may be represented as regular expression patterns. For example, a defect signature-based regular expression pattern may query for records including a text string (e.g., "test unit ready time out"). A default signature-based query may be generated that includes parameters of a system serial number 123456, the regular expression pattern "test unit ready time out," a frequency parameter of two events per hour, and a time parameter of the last two months. The query may be performed against structured records of system diagnostic data (e.g., kernel log data) from the previous two months for system(s) associated with the serial number 123456. In the event that any of the structure records of system diagnostic data (e.g., kernel log data) include the text/string "test unit ready time

out,” execution of the query would generate a positive result. If the structured records of system diagnostic data do not include the text/string “test unit ready time out,” execution of the query may not generate a positive result.

In some embodiments, defect signature-based query may be based on a health signature. A health signature may include information associated with normal or optimal operation of a system. A health signature-based query may be performed against system diagnostic data to identify conditions (e.g., CPU usage, temperature, input/output rates) that are not compliant with a health signature defining optimal or normal system operation. A health signature-based query may be used to diagnose or identify system defects or conditions not yet documented as defects (e.g., bugs or bug signatures). A health signature-based query may be used to identify outlier system(s) (e.g., system(s) with characteristics that deviate from optimal/normal operation).

At **210**, it is determined whether the query yields a positive result (e.g., a match, correlation, hit). In some cases, a default signature-based query against system diagnostic data may not yield a positive result, and the process would terminate. In this case, an indication may be output that no defects were discovered as a result of the default signature-based query. An output may also identify one or more structured records (or portions thereof) that resemble, but do not fully match, the default signature-based query.

Alternatively, a default signature-based query against system diagnostic data may result in a match, correlation, or positive result. For example, one or more of the structured records of system diagnostic data may include information matching, associated with, or corresponding to the default signature-based query parameters. And, execution of a defect signature-based query may identify defect symptoms associated with a defect signature in the structured records of system diagnostic data.

At **220**, upon a determination that the defect is associated with the system, diagnostic engine and/or a device associated therewith, may output defect identification information. The defect identification information may be used by engineers or programmers to gain deeper understanding regarding the defect. In some embodiments, defect identification information may include a portion of the system diagnostic data that includes symptoms of the defect. By way of example, if a default signature-based query identifies an element of a defect signature (e.g., a defect symptom) in a line of kernel log data from a system, the line of kernel log data or an identifier thereof may be output. In some embodiments, the defect identification information output may include a defect identifier (e.g., a serial number associated with the defect or defect signature). Alternatively, defect identification information may include a summary of the defect. A defect summary may include background information regarding a defect associated with the defect signature. The defect summary may include hardware/software versions associated with the defect, a number of occurrences of defect symptoms (e.g., in the system or a group of systems), time(s) of occurrences of defect symptoms, root cause(s) associated with a defect, and other information associated with the defect.

According to some embodiments, it may be determined that a defect signature is associated with a system. Whether a defect signature is associated with a system may be determined based at least in part on whether the results of the defect signature-based query of step **210** yields a positive result. A defect signature-based query against system diagnostic data producing a positive result (e.g., a match) may indicate that symptoms associated with a defect are present in the system. For example, if a defect signature-based query of

kernel log data identifies an instance of a defect symptom (e.g., a string representing the symptom), it may indicate that a defect signature including the defect symptom is associated with the system.

According to some embodiments, it may be determined that a defect signature is not associated with a system even though the results of the defect signature-based query of step **210** yields a positive result. For example, a defect signature-based query may in some instances yield a false positive. In this case, the defect signature (e.g., defect definition) may be modified, redefined, and/or altered to add or remove parameters. The defect signature may, for example, be modified to add additional qualifiers (e.g., defect symptoms and/or other parameters).

In some embodiments, it may be determined that the system is not compliant with a health signature. For example, if a health signature-based query of system diagnostic data indicates that performance of a system component is sub-optimal, it may be determined that the system is not compliant with a health signature associated with that system component.

FIG. 3 is a flowchart illustrating embodiments of a process for generating defect signatures. In some embodiments, the process is performed by system **100** depicted in FIG. 1. At **300**, defect symptom data is received (e.g., from an engineer, programmer, or other user) at, for example, system **100** in FIG. 1.

At **310**, the defect symptom data may be processed (e.g., parsed and otherwise processed) into one or more defect symptom records (e.g., collectively defect symptom data). The defect symptom records may be stored in a data structure (e.g., a table, index, array, data object) stored in a database.

At **320**, it is determined whether the received defect symptoms match an existing defect signature. For example, a query based on the received defect symptom(s) may be performed against one or more defect signatures (e.g., stored in structured defect signature records). In the event that defect symptom(s) match one or more of the defect signatures, the process may proceed to step **340** (discussed below). In the event that defect symptom(s) do not match one or more of the defect signatures, a new defect signature may be generated.

At **330**, a defect signature **330** is generated based at least in part on one or more defect symptom records (e.g., defect symptom data). For example, the defect signature may be generated to include the one or more defect symptom records.

At **340**, a defect signature-based query is generated based at least in part on the defect signature. The defect signature-based query may be based on any combination of defect-related parameters (e.g., defect symptoms), logical operators, techniques, and/or functions.

FIG. 4 is a flowchart illustrating embodiments of a process for processing system diagnostic data. In some embodiments, the process is performed by system **100** depicted in FIG. 1. At **400**, system diagnostic data (e.g., system log data, configuration data, registry data, telemetry data, autosupport data) are periodically received at system (e.g., once a day).

At **410**, diagnostic data may be parsed and/or otherwise processed by the diagnostic data server to generate one or more structured records. In one example, the system diagnostic data may include a wide-variety of information related to the system (e.g., including extraneous information), and this information may be parsed and filtered to extract relevant system diagnostic data. This data may be converted into structured records of system diagnostic data.

At **420**, the structured records (e.g., structured diagnostic records) may be stored in a database of system diagnostic data that includes diagnostic data from one or more systems (e.g., a group of systems).

7

FIG. 5 is a diagram depicting an example interface for defect diagnosis. A defect diagnosis interface 500 allows a user to enter defect symptom data, defect filters, search information, and other data. For example, a user may enter a date range into symptom data entry fields 510, select defect filters 520 (e.g., domain, component, defect ID and other filters (not shown)), and/or enter additional criteria (e.g., serial number of a system) into search data entry fields 530. A defect signature-based query may be generated based on the entered information and the query may be executed upon receipt of input from the scan symptom data button 540. The interface 500 is one example of an interface; the present disclosure, however, encompasses a wide-variety of interface configurations.

FIG. 6 is a diagram depicting an example interface for displaying defect diagnosis results. A defect diagnosis display interface 600 may output defect identification information including a defect identifier 610 (e.g., a bug ID), data identifying a system 620, a defect symptom summary 630 (e.g., including a timestamp, hardware/software version(s), a portion of the system diagnostic data in which the defect symptoms are identified), and other information.

Although the foregoing embodiments have been described in some detail for purposes of clarity of understanding, the invention is not limited to the details provided. There are many alternative ways of implementing the invention. The disclosed embodiments are illustrative and not restrictive.

What is claimed is:

1. A method of defect diagnosis, comprising:
  - performing, by a processor, a defect signature-based query against system diagnostic data stored in one or more structured records;
  - determining that a defect signature stored in memory is associated with a system based at least in part on the query;
  - determining whether the defect signature association results in a false positive; and
  - in response to the determined false positive, modifying the defect signature by adding additional qualifiers and storing the modified defect signature.
2. The method of claim 1, further comprising:
  - prior to performing the defect signature-based query, receiving diagnostic data from the system; and
  - generating system diagnostic data stored in one or more structured records based at least in part on the diagnostic data, wherein the structured records are generated by processing the received diagnostic data.
3. The method of claim 1, further comprising:
  - generating the defect signature-based query based at least in part on the defect signature and one or more parameters associated with the system.
4. The method of claim 3, wherein the parameters include one or more of a number of occurrences of an event, a frequency of occurrence, a code level associated with a defect symptom, a period of time, and a product version.
5. The method of claim 3, wherein the defect signature comprises one or more defect symptoms indicative of a defect.
6. The method of claim 1, wherein the defect signature-based query includes one or more logical operators.
7. The method of claim 1, wherein the defect signature-based query includes a regular expression.
8. The method of claim 1, further comprising outputting remediation information generated based at least in part on the defect signature and the system diagnostic data.
9. The method of claim 1, further comprising outputting defect identification information associated with one or more

8

of a representation of system diagnostic data that includes the defect, a defect identifier associated with the defect signature, and a defect symptom summary.

10. The method of claim 1, wherein the defect signature comprises a bug signature.

11. The method of claim 1, wherein the determining that a defect signature stored in memory is associated with a system based at least in part on the query step comprises determining that the system is not compliant with a health signature.

12. The method of claim 1, wherein the determining that a defect signature stored in memory is associated with a system based at least in part on the query step comprises determining that the defect signature is not associated with the system based at least in part on the query.

13. A system for defect diagnosis, comprising:

a processor; and

a memory coupled with the processor, wherein the memory is configured to provide the processor with instructions which when executed cause the processor to:

perform a defect signature-based query against system diagnostic data stored in one or more structured records;

determine that a defect signature is associated with a serviced system based at least in part on the query;

determine whether the defect signature association results in a false positive; and

in response to the determined false positive, modifying the defect signature by adding additional qualifiers and storing the modified defect signature.

14. The system of claim 13, wherein the memory is further configured to provide the processor with instructions which when executed cause the processor to:

prior to performing the defect signature-based query, receive diagnostic data from the serviced system; and

generate system diagnostic data stored in one or more structured records based at least in part on the diagnostic data, wherein the structured records are generated by processing the received diagnostic data.

15. The system of claim 13, wherein the memory is further configured to provide the processor with instructions which when executed cause the processor to generate the defect signature-based query based at least in part on the defect signature and one or more parameters associated with the serviced system.

16. The system of claim 15, wherein the defect signature comprises one or more defect symptoms indicative of a defect.

17. The system of claim 15, wherein the parameters include one or more of a number of occurrences of an event, a frequency of occurrence, a code level associated with a defect symptom, a period of time, and a product version.

18. The system of claim 13, wherein the memory is further configured to provide the processor with instructions which when executed cause the processor to:

output defect identification information associated with one or more of a representation of system diagnostic data that includes the defect, a defect identifier associated with the defect signature, and a defect symptom summary.

19. A computer program product for defect diagnosis, the computer program product being embodied in a tangible non-transitory computer readable storage medium and comprising computer instructions for:

performing a defect signature-based query against system diagnostic data stored in one or more structured records;

determining that a defect signature is associated with a system based at least in part on the query;

determining whether the defect signature association results in a false positive; and  
in response to the determined false positive, modifying the defect signature by adding additional qualifiers and storing the modified defect signature.

5

**20.** The computer program product of claim **19**, further comprising:

prior to performing the defect signature-based query, receiving diagnostic data from the system; and  
generating system diagnostic data stored in one or more structured records based at least in part on the diagnostic data, wherein the structured records are generated by processing the received diagnostic data.

10

\* \* \* \* \*